

IN THE UNITED STATES  
PATENT AND TRADEMARK OFFICE

In re the Application of:

Inventor: Christopher J. Chase et al Case: IDS 2000-0660CIP

Serial No.: |  
| To be assigned  
Filing Date: |

Title: Technique for Ethernet Access to Packet-Based Services

To: THE COMMISSIONER OF PATENTS AND TRADEMARKS  
WASHINGTON, D.C. 20231

Sir:

**PRELIMINARY AMENDMENT**

Prior to taking any action on the instant continuation-in-part application, kindly enter the following amendments to the specification, claims and abstract:

**IN THE SPECIFICATION:**

Page 1, lines 3-4

This invention relates to a technique enabling access to packet-based services, such as IP, Frame Relay, and ATM, through an Ethernet protocol network.

Page 1, lines 13-20

As an alternative to private line access, communications service providers such as AT&T also offer virtual circuit access allowing several customers to logically share a single circuit, thus reducing costs. Such shared circuits, typically referred to as Permanent Virtual Circuits, allow communication service providers to guarantee customer traffic flows that are distinguishable from each other, are secure, and allow customers to enjoy different service features. An example of such a technique for offering such shared service is disclosed in U.S. Patent 6,081,524, assigned to AT&T.

Page 2, lines 8-20

Briefly, in accordance with a preferred embodiment, a method is provided for routing data in an Ethernet protocol network having a plurality of platforms, each serving one or more customers. A first platform receives at least one frame from a sending site (e.g., a first customer's premises) that is destined for a receiving site (e.g., another premises belonging to the same or a different customer.) After receiving the frame, the first platform overwrites a portion of the frame with a **customer descriptor** that specifically identifies the sending customer. In practice, the first platform may overwrite a Virtual Local Area Network (VLAN) field that is typically employed by the sending customer for the purposes of routing data among various VLANs at the sending premises. Rather than overwrite the VLAN field in the frame, the first platform could overwrite another field, such as the source address field, or could even employ a "shim" header containing the customer descriptor. All further use of the term **customer descriptor** implies that any of the above or similar techniques could be used.

Page 2, lines 21-28, page 3, lines 1-2

After overwriting the frame with the customer descriptor, the sending platform routes the frame onto the MAN for routing among the other platforms, thereby sharing trunk bandwidth and other resources, but logically distinct from other customers' traffic with different customer descriptors. A destination address in the frame directs the frame to its corresponding endpoint. Upon receipt of the frame, the receiving platform uses the customer descriptor to segregate the frame for delivery to the proper destination, especially in the event where different customers served by the same platform use overlapping addressing plans. Thus, the customer descriptor in each frame advantageously enables the receiving platform to distinguish between different customers served by that platform.

Page 3, lines 3-6

For traffic with a destination beyond the MAN, this method provides a convenient and efficient way to map the customer descriptor to similar identifiers in a Wide Area Network (WAN), such as a Permanent Virtual Circuit (PVC), a Virtual Private Network (VPN), or an MPLS Label Switched Circuit.

Page 3, lines 7-13

Overwriting each frame with the customer descriptor thus affords the ability to logically segregate traffic on the Ethernet MAN to provide Virtual Private Network (VPN) services of the type offered only on more expensive Frame Relay and ATM networks. Moreover, the customer descriptor used to tag each frame can advantageously include Quality of Service (QoS) information, allowing the sender to specify different QoS levels for different traffic types, based on the Service Level Agreement (SLA) between the customer and the communications service provider.

Page 3, lines 15-17

FIGURE 1 depicts an Ethernet protocol Metropolitan Area Network (MAN) in which each frame is tagged with a customer descriptor in its VLAN field in accordance with the present principles;

Page 3, lines 18-19

FIGURE 2 illustrates a sample frame for transmission over the network of FIG. 1;

Page 3, lines s 24-26

FIGURE 5 illustrates a portion of a MAN showing the manner in which frames are mapped to different Permanent Virtual Circuits by an ATM switch;

page 3, lines 27-29

FIGURE 6 illustrates a portion of a MAN showing the manner in which frames are mapped into different Multi-Protocol Label Switching (MPLS) tunnels;

Page 4, lines 1-2

FIGURE 7 illustrates a portion of a MAN showing the manner in which frames are mapped into different service networks;

FIGURE 8 illustrates a portion of a prior art Ethernet protocol network in which the VLAN on an incoming frame received at an ingress port of a switch extends directly to frame output by the switch at an egress port; and

FIGURE 9 illustrates a portion of an alternate preferred embodiment of the invention in which a VLAN tag on a frame received at an ingress port of a switch is

mapped to a second tag that is unique to an egress port of the switch which outputs the frame.

Page 4, lines 4-10

FIGURE 1 depicts an Ethernet Protocol Metropolitan Area Network (MAN) 10 comprised of a plurality of Multi-Service Platforms (MSPs) 12<sub>1</sub>-12<sub>n</sub> where *n* is an integer, each MSP taking the form of an Ethernet switch or the like. In the illustrated embodiment *n*=4, although the network 10 could include a smaller or larger number of MSPs. A fiber ring or SONET ring infrastructure 14 connects the platforms 12<sub>1</sub>-12<sub>4</sub> in daisy-chain fashion allowing each MSP to statistically multiplex information onto, and to statistically de-multiplex information off the ring infrastructure 14.

Page 4, lines 11-21

Each of MSPs 12<sub>1</sub>-12<sub>3</sub> serves at least one, and in some instances, a plurality of premises 16 belonging to one or more customers. In the illustrated embodiment of FIG. 1, the MSP 12<sub>1</sub> serves a single customer premises 16<sub>1</sub> belonging to customer 1 whereas, the MSP 12<sub>2</sub> serves premises 16<sub>2</sub>, and 16<sub>3</sub> belonging to customers 2 and 3, respectively. The MSP 12<sub>3</sub> serves a single premises 16<sub>4</sub> that belongs to customer 3. The MSPs 12<sub>1</sub>-12<sub>3</sub> are linked to their corresponding premises via 10, 100 or 1000 MB links 19. The MSP 12<sub>4</sub> bears the legend “CO MSP” because it serves as a central office to route traffic from the MAN 10 to a Provider Edge Router (PER) 18 for delivery to other networks, such as Frame Relay, ATM, MPLS networks or the Internet as discussed hereinafter. By the same token, the PER 18 can route traffic from such other networks onto the MAN 10 via the CO MSP 12<sub>4</sub>.

Page 4, lines 20-29

The traffic routed onto and off of the MAN 10 by each MSP takes the form of one or more frames 20 depicted in FIG. 2. Heretofore, traffic routed onto the MAN 10 from a particular customer's premises was combined with other customers' traffic with no logical separation, thus raising security concerns. Moreover, since all customers' traffic share the same bandwidth, difficulties existed in prior art Ethernet MANs in regulating the traffic from each customer's premises, and in affording different customers different Quality of Service (QoS) levels in accordance with individual Service Level Agreements.

Page 5, lines 1-8

These difficulties are overcome in accordance with the present principles by "tagging" each frame 20 routed onto the MAN 10 at a particular MSP, say MSP 12<sub>3</sub>, with a customer descriptor 22' (best seen in FIG. 2) that identifies the customer sending that frame. As discussed in greater detail below, each MSP receiving a frame 20 on the fiber ring infrastructure 14 uses the customer descriptor 22' in that frame to maintain distinct routing and addressing tables that are assigned to each customer served by that MSP. This permits each customer to use its own addressing plan without fear of overlap with other customers, as the customers are all maintained as logically separate.

Page 5, lines 9-15

FIGURE 2 depicts the structure of an exemplary Ethernet protocol frame 20 specified by Ethernet Standard 802.1Q. Among the blocks of bytes within each frame 20 is a Virtual Local Area Network (VLAN) Identifier 22 comprised of sixteen bits. In practice, the VLAN identifier 22, in conjunction with a VLAN flag 23 within the frame, facilitates routing of the frame within a customer's premises to a particular VLAN. However, the VLAN identifier 22 has no influence on routing of the frame 20 after receipt at a MSP.

Page 5, lines 16-28

In accordance with the present principles, the prior disadvantages associated with conventional Ethernet networks, namely the lack of security and inability to regulate QoS levels, are overcome by overwriting the VLAN identifier 22 in each frame 20 with the customer descriptor maintained by the service provider. Overwriting the VLAN identifier 22 of FIG. 2 of each frame 20 with the customer descriptor 22' serves to "tag" that frame with the identity of its sending customer, thus affording each MSP in the MAN 10 the ability to route those frames only among the premises belonging to that same sending customer. Such tagging affords the operator of the MAN 10 the ability to provide security in connection with frames transmitted across the network, since frames with customer ID A would not be delivered to any premises of customer with ID B. As an example, two or more customers served by a single MSP may use overlapping IP addressing schemes. In the absence of any other identifier, the MSP receiving frames with overlapping IP addresses lacks the ability to assure accurate delivery.

Page 5, lines 29-31, Page 6, lines 1-3

In the illustrated embodiment depicted in FIG. 2, each MSP of Fig. 1 tags each outgoing frame 20 by overwriting the VLAN identifier 22 with the customer descriptor 22'. However, tagging could occur in other ways, rather than overwriting the VLAN identifier 22. For example, the source address block 25 within the frame 20 could be overwritten with the customer descriptor 22'. Alternatively, the data field 21 could include a shim header comprising the customer descriptor 22'.

Page 6, lines 4-15

The tagging of each frame 20 with the customer descriptor 22' affords several distinct advantages in connection with routing of the frames through the MAN 10. First, as discussed above, the tagging affords each recipient MSP the ability to distinguish traffic destined for customers with overlapping address schemes, and thus allows for

segregating traffic on the MAN 10. Further, tagging each frame 20 with the customer descriptor 22' enables mapping of the frames from a MAN 100 depicted in FIG 3 to corresponding one of a plurality of customer Virtual Private Networks 26<sub>1</sub>-26<sub>3</sub> within an MPLS network 28. As seen in FIG. 3, an MSP 120<sub>2</sub> within the MAN 100 receives traffic from premises 160<sub>1</sub>, 160<sub>2</sub>, and 160<sub>3</sub> belonging to customer 1, customer 2 and customer 3, respectively, which enjoy separate physical links to the MSP. Upon receipt of each frame from a particular customer, the MSP 120<sub>2</sub> overwrites that frame with the customer descriptor 22' corresponding to the sending customer.

Page 6, lines 16-22

After tagging each frame, the MSP 120<sub>2</sub> statistically multiplexes the frames onto the fiber ring infrastructure 14 for transmission to a CO MSP 120<sub>4</sub> for receipt at a destination PER 180 that serves the MPLS network 28 within which are customer Virtual Private Networks 26<sub>1</sub>-26<sub>3</sub>. Using the customer descriptor 22' in each frame, the PER 180 maps the frame to the corresponding VPN identifier associated with a particular one of customer Virtual Private Networks 26<sub>1</sub>-26<sub>3</sub> to properly route each frame to its intended destination.

Page 6, lines 23-31, page 7, lines 1-2

The tagging scheme of the present invention also affords the ability to route frames with different QoS levels within a MAN 1000 depicted in FIG 4. As seen in FIG. 4, an MSP 1200<sub>2</sub> within the MAN 1000 receives traffic from premises 1600<sub>2</sub>, and 1600<sub>3</sub> belonging to customer 2 and customer 3, respectively, which enjoy separate physical links to the MSP, allowing each to send frames into the MAN. In the illustrated embodiment of FIG. 4, the frames originating from the premise 1600<sub>2</sub> may contain either voice or data and have a corresponding QoS level associated with each type of frame. Upon receiving such frames, the MSP 1200<sub>2</sub> overwrites the frame with the customer descriptor 22' corresponding to the particular customer sending the frame. The

customer descriptor 22' will not only contain the identity of the sending customer, but the corresponding QoS level associated with that frame.

page 7, lines 3-10

After tagging each frame, the MSP 1200<sub>2</sub> statistically multiplexes the frames onto the fiber ring infrastructure 14 for transmission to a CO MSP 1200<sub>4</sub> for receipt at a destination PER 1800 that serves an MPLS network 280 within which are customer Virtual Private Networks 260<sub>2</sub> and 260<sub>3</sub>. Using the customer descriptor 22' in each frame, the PER 1800 of FIG. 4 maps the frame to the corresponding customer VPN to properly route each frame to its intended customer VPN. Further, the PER 1800 of FIG. 4 also maps the QoS level specified in the customer descriptor in the frame to assure that the appropriate quality of service level is applied to the particular frame.

Page 7, lines 11-27

In the above-described embodiments, the frames of customer traffic have been assumed to comprise IP packets that terminate on a router (i.e., Provider Edge Routers 18, 180 and 1800) and that the VPNs employ MPLS-BGP protocols. However, some customers may require multi-protocol support, or may otherwise require conventional PVCs so that the traffic streams must be mapped into Frame Relay or ATM PVCs as depicted in FIG. 5, which illustrates a portion of a MAN 10000 having a CO MSP12000<sub>4</sub> serving an ATM switch 30 that receives traffic from the MAN. As seen in FIG. 5, each of premises 16000<sub>1</sub>, 16000<sub>2</sub> and 16000<sub>3</sub> belonging to customer 1, customer 2 and customer 3, respectively, may send frames for receipt at MSP 12000<sub>2</sub> in the MAN 10000. The MSP 12000<sub>2</sub> tags each frame with the corresponding customer descriptor prior to statistically multiplexing the data for transmission on the fiber ring infrastructure 14 to the CO MSP 12000<sub>4</sub> for receipt at the ATM switch 30. The ATM switch 30 then maps each frame to the appropriate PVC in accordance with the customer descriptor 22' in the frame in a manner similar to the mapping described with respect to FIG. 3. Thus, the

ATM switch 30 could map the frame to one of Frame Relay recipients' 32<sub>1</sub>, 32<sub>2</sub>, or 32<sub>3</sub>, ATM recipients 32<sub>4</sub> or 32<sub>5</sub> or IMA (Inverse Multiplexing over ATM) recipient 32<sub>6</sub>.

Page 7, lines 28-31, Page 8, lines 1-6

FIG. 6 depicts a portion of a MAN network 100000 that routes frames onto separate MPLS tunnels 40<sub>1</sub>-40<sub>3</sub> (each emulating a private line 32 in an MPLS network 28000) in accordance with the customer descriptor 22' written into each frame by a MSP 120000<sub>2</sub> in the MAN. Each of customer premises 160000<sub>1</sub>, 160000<sub>2</sub> and 160000<sub>3</sub> depicted in FIG. 6 sends information frames for receipt at MSP 120000<sub>2</sub>. The MSP 120000<sub>2</sub> tags each frame with the customer descriptor prior to statistically multiplexing the data for transmission on the fiber ring infrastructure 14 for delivery to a CO MSP 120000<sub>4</sub> that serves a PER 18000. The PER 18000 translates (maps) the customer descriptors written onto the frames by the MSP 120000<sub>2</sub> into the MPLS tunnels 40<sub>1</sub>-40<sub>3</sub> to enable the PER to route the traffic to the intended customer.

Page 8, lines 7-17

FIG. 7 depicts a portion of a MAN network 1000000 for routing traffic (i.e., frames) onto separate networks in accordance with the customer descriptor written into each the frame by a MSP 120000<sub>2</sub> in the MAN. Each of customer premises 1600000<sub>2</sub> and 1600000<sub>3</sub> depicted in FIG. 7 sends frames for receipt by the MSP 1200000<sub>2</sub>. The MSP 1200000<sub>2</sub> tags each frame with the customer descriptor 22' prior to statistically multiplexing the data for transmission on the fiber ring infrastructure 14 for delivery to a CO MSP 1200000<sub>4</sub> that serves a PER 180000. In accordance with the customer descriptor, the PER 180000 of FIG. 7 routes traffic to a particular one of several different networks, e.g., an Intranet VPN 42<sub>1</sub>, a voice network 42<sub>2</sub> and the Internet 42<sub>3</sub>, in accordance with the customer descriptor 22' written onto the frame by the MSP 1200000<sub>2</sub>.

Page 8, before line 18, insert--

Referring to FIG. 8, a prior art Ethernet switch 20000000 receives Ethernet frames at one of a plurality of input (ingress) ports, exemplified by ports 22000000, 24000000, and 26000000, from one a corresponding one of networks 28000000, 30000000 and 32000000, respectively. The frames are destined for an endpoint (not shown) served by a Wide Area Network (WAN) 36000000 linked to an egress port 40000000 of the switch 20000000 by an Ethernet trunk 38000000. Each Ethernet frame received at one of the ingress switch ports 22000000, 24000000, and 26000000 carries a tag, which in accordance with the IEEE 802.1Q Standard, identifies the Virtual Local Area Network (VLAN) that originated the frame. Thus, for example, a frame originated at network 32000000 associated with a VLAN having an Identification Designation (ID) of 5 will carry a tag with the corresponding VLAN ID. The VLAN address is twelve bits, offering the ability designate as many as 4096 separate VLANs.

A VLAN domain extends across any set of connected Ethernet switches, and therefore the address space of 4096 individual VLANs is shared across such an extended network of switches. In the past, the VLAN tag associated with an incoming Ethernet frame received at one of the ingress switch ports will extend directly to the egress switch port. Hence, the VLAN tag of an Ethernet frame received at the ingress port 26000000 extends directly to the egress port 40000000 on which the switch outputs the frame. The direct extension of the VLAN tag between the Ethernet switch ingress and egress ports increases the difficulty in the sharing and administration of the limited VLAN address space, as it now has to be coordinated across any connected group of Ethernet networks, even if they only are connected by termination on a common WAN access switch, as shown in Figure 8. It also limits the size of a single switch in terms of VLAN capacity, being confined to 4096 VLANs on any given switch.

Referring to FIG. 9, in accordance with the present invention, the significance of the VLAN tag is localized to each physical port on the Ethernet switch 2000000, instead of being global to a network. At an ingress switch port, say port 22000000, the VLAN tag is still used to discriminate between different customer's traffic or services, but the switch 2000000 is free to re-write the tag to another value that is unique to the physical

egress port 40000000. In other words, the switch 20000000 may terminate traffic from many independent networks, each using the full 4096 VLAN address space, and internally map the traffic using a unique tuple of (Physical port, VLAN ID) to the switch output ports (only one of which is shown). This dramatically increases the scale achievable with a single switch, which is, by virtue of the mapping of tags from an ingress to egress port is now limited only by 4096 VLAN IDs on *each* physical port, rather than a total 4096 VLANs as is the case of the prior network of FIG. 8.

IN THE CLAIMS:

Cancel Claims 1-25 and insert claims 26-30

1        26.     In an Ethernet protocol network having at least one switch with plurality  
2     of ingress ports that are each adapted to receive at least one Ethernet frame that includes a  
3     tag that identifies a particular network sending the frame, and the switch having at least  
4     one egress port on which the frame is output, a method for operating said switch,  
5     comprising the step of:

6              mapping the tag in the Ethernet frame received at the one ingress port to a second  
7     tag associated with the egress port through which the switch outputs the frame; and

8              overwriting the tag in the Ethernet frame with the second tag prior to outputting  
9     the frame on the egress port.

1        27.     The method according to claim 26 wherein the tag is mapped using a  
2     unique tuple of the port and a Virtual Local Area Network (VLAN) identifier.

1        28.     The method according to claim 27 wherein the Virtual Local Area  
2     Network (VLAN) identifier has an prescribed address space and wherein each egress port  
3     can support a quantity of VLANS limited only by the prescribed address space of the  
4     VLAN identifier.

1           29.     The method according to claim 28 wherein the VLAN identifier has an  
2 address space of 4096 and wherein each egress port can support 4096 separate VLANs.

1           30.    In an Ethernet protocol network having at least one switch with plurality  
2 of ingress ports that are each adapted to receive at least one Ethernet frame that includes a  
3 Virtual Local Area Network (VLAN) ID tag that identifies a particular network sending  
4 the frame to that ingress port, and the switch has at least one egress port on which the  
5 frame is output, a method for operating said switch, comprising the step of:

6               mapping the tag in the Ethernet frame received at the one ingress port to a second  
7 tag using a unique tuple of the port and a Virtual Local Area Network (VLAN) identifier;  
8 and

9               overwriting the tag in the Ethernet frame received at the one ingress port with the  
10 second tag prior to outputting the frame on the egress port.

IN THE DRAWINGS

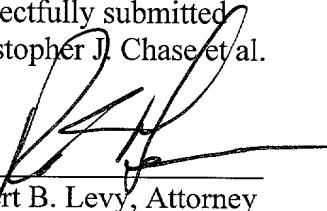
Kindly enter newly submitted FIGURES 8 and 9.

REMARKS

The foregoing amendment has been tendered in connection with the instant continuation-in-part application.

In the event that any issues remain following entry of this amendment, applicant's attorney invites the examiner to contact him (908) 221-5714 for either a personal or telephone interview if the examiner believes that such would expedite the prosecution of this application.

Respectfully submitted  
Christopher J. Chase et al.

By   
Robert B. Levy, Attorney  
Registration No. 28,234

AT&T

10/31/01

Date

MARKED-UP VERSION SHOWING CHANGES

IN THE SPECIFICATION:

Page 1, lines 3-4

This invention relates to a technique enabling access to packet-based services, such as IP, Frame Relay, and ATM, through an Ethernet [P]protocol network.

Page 1, lines 13-20

As an alternative to private line access, communications service providers such as AT&T also offer virtual circuit access allowing several customers to logically share a single circuit, thus reducing costs. Such shared circuits, typically referred to as Permanent Virtual Circuits, allow communication service providers to guarantee customer traffic flows that are distinguishable from each other, are secure, and allow customers to enjoy different service features. An example of such a technique for offering such shared service [in a Multi-Protocol Label Switching Network] is disclosed in U.S. Patent 6,081,524, assigned to AT&T.

Page 2, lines 8-20

Briefly, in accordance with a preferred embodiment, a method is provided for routing data in an Ethernet protocol network having a plurality of platforms, each serving one or more customers. A first platform receives at least one frame from a sending site (e.g., a first customer's premises) that is destined for a receiving site (e.g., another premises belonging to the same or a different customer.) After receiving the frame, the first platform overwrites a portion of the frame with a **customer descriptor** that specifically identifies the sending customer. In practice, the first platform [will] may overwrite a Virtual Local Area Network (VLAN) field that is typically employed by the sending customer for the purposes of routing data among various VLANs at the sending premises [premises]. Rather than overwrite the VLAN field in the frame, the first

platform could overwrite another field, such as the source address field, or could even employ a “shim” header containing the customer descriptor. All further use of the term **customer descriptor** implies that any of the above or similar techniques could be used.

Page 2, lines 21-28, page 3, lines 1-2

After overwriting the frame with the customer descriptor, the sending platform routes the frame onto the MAN [network] for routing among the other platforms, thereby sharing trunk bandwidth and other resources, but logically distinct from other [customer’s] customers’ traffic with different customer descriptors. A destination address in the frame directs the frame to its corresponding [receiving platform] endpoint. Upon receipt of the frame, the receiving platform uses the customer descriptor to segregate the frame for delivery to the proper destination, especially in the event where different customers served by the same platform use overlapping addressing plans. Thus, the customer descriptor in each frame advantageously enables the receiving platform to distinguish between different customers served by that platform.

Page 3, lines 3-6

For traffic with a destination beyond the MAN, this method provides a convenient and efficient way to map the customer[-]descriptor to similar identifiers in a Wide Area Network (WAN), such as a Permanent Virtual Circuit (PVC), a Virtual Private Network (VPN), or an MPLS Label Switched Circuit.

Page 3, lines 7-13

Overwriting each frame with the customer[-]descriptor thus affords the ability to logically segregate traffic on the Ethernet MAN to provide Virtual Private Network (VPN) services of the type offered only on more expensive Frame Relay and ATM networks. Moreover, the customer descriptor used to tag each frame can advantageously include Quality of Service (QoS) information, allowing the sender to specify different QoS levels for different traffic types, based on the Service Level Agreement (SLA) between the customer and the communications service provider.

Page 3, lines 15-17

FIGURE 1 depicts an Ethernet [P]rotocol Metropolitan Area Network (MAN) in which each [information] frame is tagged with a customer descriptor in its VLAN field in accordance with the present principles;

Page 3, lines 18-19

FIGURE 2 illustrates a sample [information] frame for transmission over the network of FIG. 1;

Page 3, lines 24-26

FIGURE 5 illustrates a portion of a MAN showing the manner in which [information] frames are mapped to different Permanent Virtual Circuits by an ATM switch;

Page 3, lines 27-29

FIGURE 6 illustrates a portion of a MAN showing the manner in which [information] frames are mapped into different Multi-Protocol Label Switching (MPLS) tunnels; [and]

Page 4, lines 1-2

FIGURE 7 illustrates a portion of a MAN showing the manner in which [information] frames are mapped into different service networks[.];

FIGURE 8 illustrates a portion of a prior art Ethernet protocol network in which the VLAN on an incoming frame received at an ingress port of a switch extends directly to frame output by the switch at an egress port; and

FIGURE 9 illustrates a portion of an alternate preferred embodiment of the invention in which a VLAN tag on a frame received at an ingress port of a switch is mapped to a second tag that is unique to an egress port of the switch which outputs the frame.

Page 4, lines 4-10

FIGURE 1 depicts an Ethernet Protocol Metropolitan Area Network (MAN) 10 comprised of a plurality of Multi-Service Platforms (MSPs) 12<sub>1</sub>-12<sub>n</sub> where n is an integer, each MSP taking the form of an Ethernet switch or the like. In the illustrated embodiment n=4, although the network 10 could include a smaller or larger number of MSPs. A fiber ring or SONET ring infrastructure 14 connects the platforms 12<sub>1</sub>-12<sub>4</sub> in daisy-chain fashion allowing each MSP to statistically multiplex information onto, and to statistically de-multiplex[ing] information off the ring infrastructure 14.

Page 4, lines 11-21

Each of MSPs  $12_1$ - $12_3$  serves at least one, and in some instances, a plurality of premises  $16$  belonging to one or more customers. In the illustrated embodiment of FIG. 1, the MSP  $12_1$  serves a single customer premises  $16_1$  belonging to customer 1 whereas, the MSP  $12_2$  serves premises  $16_2$ , and  $16_3$  belonging to customers 2 and 3, respectively. The MSP  $12_3$  serves a single premises  $16_4$  that belongs to customer 3. The MSPs  $12_1$ - $1[3]2_3$  are linked to their corresponding premises via 10, 100 or 1000 MB links[18] 19. The MSP  $12_4$  bears the legend “CO MSP” because it serves as a central office to route traffic from the MAN 10 to a Provider Edge Router (PER) 18 for delivery to other networks, such as Frame Relay, ATM, MPLS networks or the Internet as discussed hereinafter. By the same token, the PER 18 can route traffic from such other networks onto the MAN 10 via the CO MSP  $12_4$ .

Page 4, lines 22-29

The traffic routed onto and off of the MAN 10 by each MSP takes the form of one or more [information] frames 20 depicted in FIG. 2. Heretofore, traffic routed onto the MAN [network] 10 from a particular customer’s premises was combined with other customer[’s] traffic with no logical separation, thus raising security concerns. Moreover, since all customers’ traffic share the same bandwidth, difficulties existed in prior art Ethernet MANs in regulating the traffic from each customer’s premises, and in affording different customers different Quality of Service (QoS) levels in accordance with individual Service Level Agreements.

Page 5, lines 1-8

These difficulties are overcome in accordance with the present principles by “tagging” each frame 20 routed onto the [network] MAN 10 at a particular MSP, say MSP  $12_3$ , with a customer descriptor 22’ (best seen in FIG. 2) that identifies the customer sending that frame. As discussed in greater detail below, each MSP receiving a frame 20 on the fiber ring infrastructure 14 uses the customer descriptor 22’ in that frame to

maintain distinct routing and addressing tables that are assigned to each customer served by that MSP. This permits each customer to use [their] its own addressing plan without fear of overlap with other customers, as the[y] customers are all maintained as logically separate.

Page 5, lines 9-15

FIGURE 2 depicts the structure of an exemplary Ethernet protocol frame 20 specified by Ethernet Standard 802.1Q. Among the blocks of bytes within each frame 20 is a Virtual Local Area Network (VLAN) Identifier 22 comprised of sixteen bits. In practice, the VLAN identifier 22, in conjunction with a VLAN flag [block] 23 within the frame, facilitates routing of the frame within a customer's premises to a particular VLAN. However, the VLAN identifier 22 has no influence on routing of the frame 20 after receipt at a MSP.

Page 5, lines 16-28

In accordance with the present principles, the prior disadvantages associated with conventional Ethernet networks, namely the lack of security and inability to regulate QoS levels, are overcome by overwriting the VLAN identifier 22 in each frame 20 with the customer descriptor maintained by the service provider. Overwriting the VLAN identifier 22 of FIG. 2 of each frame 20 with the customer descriptor 22' serves to "tag" that frame with the identity of its sending customer [identity], thus affording each MSP in the [network] MAN 10 the ability to route those frames only among the premises belonging to that same sending customer. Such tagging affords the operator of the [network] MAN 10 the ability to provide security in connection with frames transmitted across the network, since frames with customer ID A would not be delivered to any premises of customer with ID B. As an example, two or more customers served by a single MSP may use overlapping IP addressing schemes. In the absence of any other identifier, the MSP receiving [such] frames with overlapping IP addresses lacks the ability to assure accurate delivery.

Page 5, lines 29-31, Page 6, lines 1-3

In the illustrated embodiment depicted in FIG. 2, each MSP of Fig. 1 tags [the] each outgoing frame 20 by overwriting the VLAN identifier 22 with the customer descriptor 22'. However, tagging could occur in other ways, rather than overwriting the VLAN identifier 22. For example, the source address block 25 within the frame 20 could be overwritten with the customer descriptor 22'. Alternatively, the data field [25] 21 could include a shim header comprising the customer descriptor 22'.

Page 6, lines 4-15

The tagging of each frame 20 with the customer descriptor 22' affords several distinct advantages in connection with routing of the frames through the MAN 10. First, as discussed above, the tagging affords each recipient MSP the ability to distinguish traffic destined for customers with overlapping address schemes, and thus allows for segregating traffic on the MAN 10. Further, tagging each frame 20 with the customer descriptor 22' enables mapping of the frames from a MAN 100 depicted in FIG 3 to corresponding one of a plurality of customer Virtual Private Networks 26<sub>1</sub>-26<sub>3</sub> within an MPLS network 28. As seen in FIG. 3, an MSP 120<sub>2</sub> within the MAN 100 receives traffic from premises 160<sub>1</sub>, 160<sub>2</sub>, and 160<sub>3</sub> belonging to customer 1, customer 2 and customer 3, respectively, which enjoy separate physical links to the MSP. Upon receipt of each frame from a particular customer, the MSP 120<sub>2</sub> overwrites that frame with the customer descriptor 22' corresponding to the sending customer.

Page 6, lines 16-22

After tagging each [data] frame, the MSP 120<sub>2</sub> statistically multiplexes the frames onto the fiber ring infrastructure 14 for transmission to a CO MSP 120<sub>4</sub> for receipt at a destination PER 180 that serves the MPLS network 28 within which are customer Virtual Private Networks 26<sub>1</sub>-26<sub>3</sub>. Using the customer descriptor 22' in each frame, the PER 180

maps the frame to the corresponding VPN identifier associated with a particular one of customer Virtual Private Networks 26<sub>1</sub>-26<sub>3</sub> to properly route each frame to its intended destination.

Page 6, lines 23-31, page 7, lines 1-2

The tagging scheme of the present invention also affords the ability to route [information] frames with different QoS levels within a MAN 1000 depicted in FIG 4. As seen in FIG. 4, an MSP 1200<sub>2</sub> within the MAN 1000 receives traffic from premises 1600<sub>2</sub>, and 1600<sub>3</sub> belonging to customer 2 and customer 3, respectively, which enjoy separate physical links to the MSP, allowing each to send [information] frames into the MAN. In the illustrated embodiment of FIG. 4, the frames originating from the premise 1600<sub>2</sub> may contain either voice or data and have a corresponding QoS level associated with each type of frame. Upon receiving such frames, the MSP 1200<sub>2</sub> overwrites the frame with the customer descriptor 22' corresponding to the particular customer sending the frame. The customer descriptor 22' will not only contain the identity of the sending customer, but the corresponding QoS level associated with that frame.

page 7, lines 3-10

After tagging each [data] frame, the MSP 1200<sub>2</sub> statistically multiplexes the frames onto the fiber ring infrastructure 14 for transmission to a CO MSP 1200<sub>4</sub> for receipt at a destination PER 1800 that serves an MPLS network 280 within which are customer Virtual Private Networks 260<sub>2</sub> and 260<sub>3</sub>. Using the customer descriptor 22' in each frame, the PER 1800 of FIG. 4 maps the frame to the corresponding customer VPN to properly route each frame to its intended customer [premises] VPN. Further, the PER 1800 of FIG. 4 also maps the QoS level specified in the customer descriptor in the frame to assure that the appropriate quality of service level is applied to the particular frame.

Page 7, lines 11-27

In the above-described embodiments, the frames of customer traffic have been assumed to comprise IP packets that terminate on a router (i.e., Provider Edge Routers 18, 180 and 1800) and that the VPNs employ MPLS-BGP protocols. However, some customers may require multi-protocol support, or may otherwise require conventional PVCs so that the traffic streams must be mapped into Frame Relay or ATM PVCs as depicted in FIG. 5, which illustrates a portion of a MAN 10000 having a CO MSP12000<sub>4</sub> serving an ATM switch 30 that receives traffic from the MAN. As seen in FIG. 5, each of premises 16000<sub>1</sub>, 16000<sub>2</sub> and 16000<sub>3</sub> belonging to customer 1, customer 2 and customer 3, respectively, [5] may [originate information] send frames for receipt at MSP 12000<sub>2</sub> in the MAN 10000. The MSP 12000<sub>2</sub> tags each frame with the corresponding customer descriptor prior to statistically multiplexing the data for transmission on the fiber ring infrastructure 14 to the CO MSP 12000<sub>4</sub> for receipt at the ATM switch 30. The ATM switch 30 then maps [the] each frame to the appropriate PVC in accordance with the customer descriptor 22' in the frame in a manner similar to the mapping described with respect to FIG. 3. Thus, the ATM switch 30 could map the frame to one of Frame Relay recipients' 32<sub>1</sub>, 32<sub>2</sub>, or 32<sub>3</sub>, ATM recipients 32<sub>4</sub> or 32<sub>5</sub> or IMA (Inverse Multiplexing over ATM) recipient 32<sub>6</sub>.

Page 7, lines 28-31, Page 8, lines 1-6

FIG. 6 depicts a portion of a MAN network 100000 that routes frames onto separate MPLS tunnels 40<sub>1</sub>-40<sub>3</sub> (each emulating a private line 32 in an MPLS network 28000) in accordance with the customer descriptor 22' written into each frame by a MSP 120000<sub>2</sub> in the MAN. Each of customer premises 160000<sub>1</sub>, 160000<sub>2</sub> and 160000<sub>3</sub> depicted in FIG. 6 [originate] sends information frames for receipt at MSP 120000<sub>2</sub>. The MSP 120000<sub>2</sub> tags each frame with the customer descriptor prior to statistically multiplexing the data for transmission on the fiber ring infrastructure 14 for delivery to a CO MSP 120000<sub>4</sub> that serves a PER 18000. The PER 18000 translates (maps) the

customer descriptors written onto the frames by the MSP 120000<sub>2</sub> into the MPLS tunnels 40<sub>1</sub>-40<sub>3</sub> to enable the PER to route the traffic to the intended customer.

Page 8, lines 7-17

FIG. 7 depicts a portion of a MAN network 1000000 for routing traffic (i.e., [information] frames) onto separate networks in accordance with the customer descriptor written into each the frame by a MSP 120000<sub>2</sub> in the MAN. Each of customer premises 1600000<sub>2</sub> and 16000003 depicted in FIG. 7 [originates information] sends frames for receipt by the MSP 120000<sub>2</sub>. The MSP 1200000<sub>2</sub> tags each frame with the customer descriptor 22' prior to statistically multiplexing the data for transmission on the fiber ring infrastructure 14 for delivery to a CO MSP 1200000<sub>4</sub> that serves a PER 180000. In accordance with the customer descriptor, the PER 1800000 of FIG. 7 routes traffic to a particular one of several different networks, e.g., an Intranet VPN 42<sub>1</sub>, a voice network 42<sub>2</sub> and the Internet 42<sub>3</sub>, in accordance with the customer descriptor 22' written onto the frame by the MSP 1200000<sub>2</sub>.

Page 8, before line 18 insert

Referring to FIG. 8, a prior art Ethernet switch 20000000 receives Ethernet frames at one of a plurality of input (ingress) ports, exemplified by ports 22000000, 24000000, and 26000000, from one a corresponding one of networks 28000000, 30000000 and 32000000, respectively. The frames are destined for an endpoint (not shown) served by a Wide Area Network (WAN) 36000000 linked to an egress port 40000000 of the switch 20000000 by an Ethernet trunk 38000000. Each Ethernet frame received at one of the ingress switch ports 22000000, 24000000, and 26000000 carries a tag, which in accordance with the IEEE 802.1Q Standard, identifies the Virtual Local Area Network (VLAN) that originated the frame. Thus, for example, a frame originated at network 32000000 associated with a VLAN having an Identification Designation (ID) of 5 will carry a tag with the corresponding VLAN ID. The VLAN address is twelve bits, offering the ability designate as many as 4096 separate VLANs.

A VLAN domain extends across any set of connected Ethernet switches, and therefore the address space of 4096 individual VLANs is shared across such an extended network of switches. In the past, the VLAN tag associated with an incoming Ethernet frame received at one of the ingress switch ports will extend directly to the egress switch port. Hence, the VLAN tag of an Ethernet frame received at the ingress port 26000000 extends directly to the egress port 40000000 on which the switch outputs the frame. The direct extension of the VLAN tag between the Ethernet switch ingress and egress ports increases the difficulty in the sharing and administration of the limited VLAN address space, as it now has to be coordinated across any connected group of Ethernet networks, even if they only are connected by termination on a common WAN access switch, as shown in Figure 8. It also limits the size of a single switch in terms of VLAN capacity, being confined to 4096 VLANs on any given switch.

Referring to FIG. 9, in accordance with the present invention, the significance of the VLAN tag is localized to each physical port on the Ethernet switch 2000000, instead of being global to a network. At an ingress switch port, say port 22000000, the VLAN tag is still used to discriminate between different customer's traffic or services, but the switch 2000000 is free to re-write the tag to another value that is unique to the physical egress port 40000000. In other words, the switch 2000000 may terminate traffic from many independent networks, each using the full 4096 VLAN address space, and internally map the traffic using a unique tuple of (Physical port, VLAN ID) to the switch output ports (only one of which is shown). This dramatically increases the scale achievable with a single switch, which is, by virtue of the mapping of tags from an ingress to egress port is now limited only by 4096 VLAN IDs on *each* physical port, rather than a total 4096 VLANs as is the case of the prior network of FIG. 8.

#### IN THE CLAIMS:

Cancel Claims 1-25 and insert claims 26-30

- 1       --26. In an Ethernet protocol network having at least one switch with plurality
- 2       of ingress ports that are each adapted to receive at least one Ethernet frame that includes a

3 tag that identifies a particular network sending the frame, and the switch having at least  
4 one egress port on which the frame is output, a method for operating said switch,  
5 comprising the step of:

6 mapping the tag in the Ethernet frame received at the one ingress port to a second  
7 tag associated with the egress port through which the switch outputs the frame; and  
8 overwriting the tag in the Ethernet frame with the second tag prior to outputting  
9 the frame on the egress port.

27. The method according to claim 26 wherein the tag is mapped using a unique tuple of the port and a Virtual Local Area Network (VLAN) identifier.

1 28. The method according to claim 27 wherein the Virtual Local Area  
2 Network (VLAN) identifier has an prescribed address space and wherein each egress port  
3 can support a quantity of VLANS limited only by the prescribed address space of the  
4 VLAN identifier.

1 29. The method according to claim 28 wherein the VLAN identifier has an  
2 address space of 4096 and wherein each egress port can support 4096 separate VLANs.

1 30. In an Ethernet protocol network having at least one switch with plurality  
2 of ingress ports that are each adapted to receive at least one Ethernet frame that includes a  
3 Virtual Local Area Network (VLAN) ID tag that identifies a particular network sending  
4 the frame to that ingress port, and the switch has at least one egress port on which the  
5 frame is output, a method for operating said switch, comprising the step of:

6 mapping the tag in the Ethernet frame received at the one ingress port to a second  
7 tag using a unique tuple of the port and a Virtual Local Area Network (VLAN) identifier;  
8 and

9 overwriting the tag in the Ethernet frame received at the one ingress port with the  
10 second tag prior to outputting the frame on the egress port. --

## ABSTRACT OF THE DISCLOSURE

An Ethernet Metropolitan Area Network (10) provides connectivity to one or more customer premises (16<sub>1</sub>, 16<sub>2</sub>, 16<sub>3</sub>) to packet-based services, such as ATM, Frame Relay, or IP, while advantageously providing a mechanism for assuring security and regulation of customer traffic. Upon receipt of each customer-generated information frame (20), an ingress Multi-Service Platform (MSP) (12<sub>2</sub>) “tags” the frame with a customer descriptor (22') that specifically identifies the recipient customer. In practice, the MSP tags each frame by overwriting the Virtual Local Area Network (VLAN) identifier (22) with the customer descriptor. Using the customer descriptor in each frame, a recipient Provider Edge Router (PER) (18) or ATM switch can map the information as appropriate to direct the information to the specific customer at its receiving site. In addition, the customer descriptor (22') may also include Quality of Service (QoS) information, allowing the recipient Provider Edge Router (PER) (18) or ATM switch to afford the appropriate QoS level accordingly. Each Ethernet switch may advantageously overwrite the VLAN identifier at an incoming port with a second tag associated with an egress port to increase the scale associated with single switch.

Fig. 8

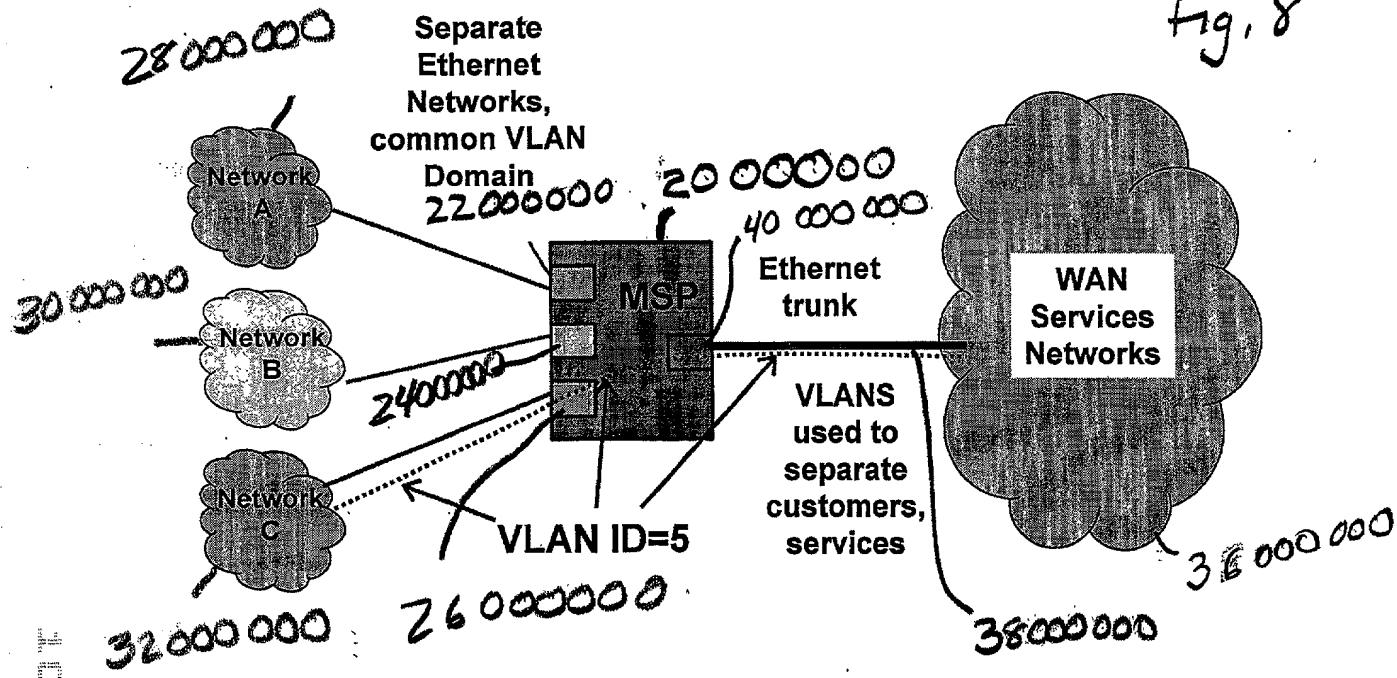


Fig. 9

